arXiv:1504.03509v2 [cs.LG] 4 Sep 2015

Regret vs. Communication: Distributed Stochastic Multi-Armed Bandits and Beyond

Shuang Liu
Shanghai Jiao Tong University
liushuang93006@gmail.com

Cheng Chen
Shanghai Jiao Tong University
jackchen1990@gmail.com

Zhihua Zhang
Shanghai Jiao Tong University
zhang-zh@cs.sjtu.edu.cn

Abstract

In this paper, we consider the distributed stochastic multi-armed bandit problem, where a global arm set can be accessed by multiple players independently. The players are allowed to exchange their history of observations with each other at specific points in time. We study the relationship between regret and communication. When the time horizon is known, we propose the Over-Exploration strategy, which only requires one-round communication and whose regret does not scale with the number of players. When the time horizon is unknown, we measure the frequency of communication through a new notion called the density of the communication set, and give an exact characterization of the interplay between regret and communication. Specifically, a lower bound is established and stable strategies that match the lower bound are developed. The results and analyses in this paper are specific but can be translated into more general settings.

1 Introduction

We consider the distributed stochastic multi-armed bandit (MAB) problem, where a global arm set $\mathcal{A} = [K]$ can be accessed by M players independently. Each arm $a \in \mathcal{A}$ is associated with an unknown yet fixed probability distribution ν_a that belongs to a known family $\mathcal{P}$. The process proceeds in rounds. At the beginning of each round t, each player $p \in [M]$ pulls an arm $A_{p,t} \in \mathcal{A}$ based on some *policy* and independently receives a reward $X_{p,t} \sim \nu_{A_{p,t}}$. Some rounds are *Communication rounds* at the end of which every player knows everything other players know. Note that when $M = 1$, the process reduces to a traditional MAB process.

The goal is to minimize the *regret*. We denote by μ_a the mean of the distribution ν_a and let $\mu^* = \max_a \mu_a$. The regret after T rounds is defined by

$$\mathscr{R}_T \coloneqq TM\mu^* - \mathbb{E}\left[\sum_{t=1}^{T}\sum_{p=1}^{M} X_{p,t}\right].$$

We are only interested in *consistent* policies. A policy is said to be consistent if for any $c > 0$, $\mathscr{R}_T = o(T^c)$ always holds. If we further let $\Delta_a = \mu^* - \mu_a$ and $N_T(a)$ be the number of times arm a has been pulled by all the players in the first T rounds, then it suffices to bound $N_T(a)$ for all a such that $\mu_a \neq \mu^*$ since the regret can be written as $\sum_{a\in\mathcal{A},\mu_a\neq\mu^*} \Delta_a \mathbb{E}\left[N_T(a)\right]$.

1.1 Related Work

The traditional single-player MAB problem has been studied for a long time. The establishment of the lower bounds can be traced back to [4, 11], in which some asymptotically policies based on the notion of *upper confidence bound* (UCB) are also developed. Later contributions mainly focus

on finite-time analysis of UCB-like algorithms [1, 5, 6, 13]. Recently, a Bayesian approach called Thompson sampling [15] is also proved to be asymptotically optimal [10].

Although there have been regret analyses of single-player models with *delayed feedback* [9], regret analysis of distributed MAB models remains largely unaddressed in the literature. The model used in [7] is very similar to ours, but they focused on best arm identification, which is mainly an exploration problem. In [14], a P2P-like gossip based model is proposed. However, their policy is based on the ϵ-GREEDY algorithm [1] which itself requires a lower bound on the gap between the best arm and the second-best arm. Furthermore, [14] only considered the case that the number of communication rounds grows linearly with T. There are also other distributed models in which players *compete* with each other [12] or an adversarial setting is considered [3].

2 Oblivious Bandit Policies

In this section we first propose a general framework for a large number of MAB policies. Then we show how this framework enables us to develop communication strategies independent of the bandit policy used.

Definition 1 (Oblivious Bandit Policy)**.** *Given K finite sets[1] of rewards $\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_K$ generated by K arms, an oblivious bandit policy chooses the next arm based only on these K sets.*

In order to explain how current bandit policies can be translated into this framework and further adapted into a distributed setting, we require some additional notation. Given a finite set of rewards $\mathcal{X}_a$ generated by arm a, the *empirical distribution* with respect to $\mathcal{X}_a$ is defined by

$$\hat{\nu}_{\mathcal{X}_a} \coloneqq \frac{1}{|\mathcal{X}_a|} \sum_{x \in \mathcal{X}_a} \delta(x),$$

where $\delta(\cdot)$ is the Dirac delta function. For every player p, every round t, and every arm a, we denote by $\mathcal{X}_{p,t}(a)$ the set of rewards generated by arm a that are available to player p at the end of round t and let $N_{p,t}(a) = |\mathcal{X}_{p,t}(a)|$. Clearly, if t is a communication round, then $N_t(a) = N_{p,t}(a)$ for every p and a. For simplicity, we use $\mathbb{E}[\nu]$ to denote the mean of distribution ν and denote $\mathbb{E}\left[\hat{\nu}_{\mathcal{X}_{p,t}(a)}\right]$ by $\tilde{\mu}_{p,t}(a)$. We also use $\mathcal{B}(p)$ to denote a Bernoulli distribution with parameter p, and $D_{\mathrm{KL}}(\nu_1||\nu_2)$ to denote the Kullback-Leibler (KL) divergence of ν_2 from ν_1. The following are two oblivious bandit policies that have been adapted into a distributed setting (and will be mainly discussed in this paper).

UCB adaptation for [0, 1] bounded rewards. Let $\mathcal{F}(t) = \ln\left(t \ln^3(t)\right)$. $A_{p,t} = \arg\max_a B^+_{p,t}(a)$

$$\text{where } B^+_{p,t}(a) = \tilde{\mu}_{p,t-1}(a) + \sqrt{\frac{\mathcal{F}\left(\sum_{k=1}^K N_{p,t-1}(k)\right)}{2N_{p,t-1}(a)}}.$$

KL-UCB adaptation for Bernoulli rewards. Let $\mathcal{F}(t) = \ln\left(t \ln^3(t)\right)$. $A_{p,t} = \arg\max_a B^+_{p,t}(a)$

$$\text{where } B^+_{p,t}(a) = \sup\left\{ p \in (0,1) : D_{\mathrm{KL}}\left(\hat{\nu}_{\mathcal{X}_{p,t-1}(a)}||\mathcal{B}(p)\right) \le \frac{\mathcal{F}\left(\sum_{k=1}^K N_{p,t-1}(k)\right)}{N_{p,t-1}(a)} \right\}.$$

There are many other oblivious bandit policies in the literature such as most UCB-like policies and Thompson sampling. These policies are called *oblivious* because their choice of the next arm depends only on the empirical distribution of the data and the number of data collected. They do not rely on a *timer* or other information such as the player's own previous choices or other players' previous choices. Note that there are non-oblivious policies such as the DMED policy [8].

3 Distributed Bandits: A Paradox

It is often believed that the more you know, the better you will do. Translating into bandit language, the more you communicate, the lower the regret is. However, the following example shows this is not necessarily true.

[1] Actually these sets are multisets

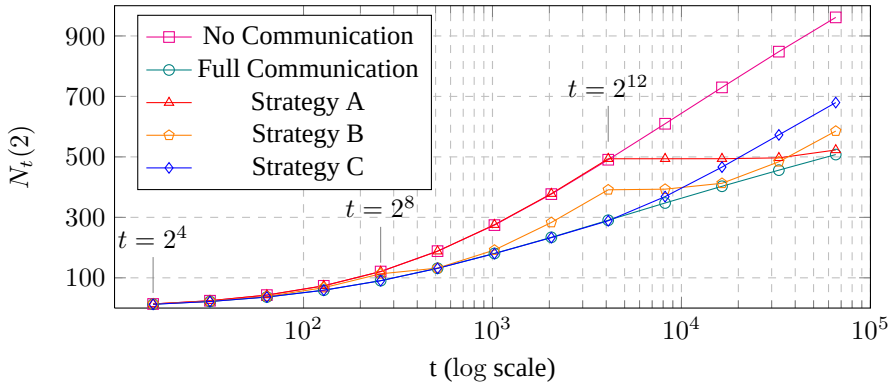

Figure 1: Adding communication rounds may increase the regret. Strategy A communicates only one time ($t = 2^{12}$). Strategy B communicates three times ($t = 2^4, 2^8, 2^{12}$). Strategy C communicates over four thousand times ($t = 1, 2, \ldots, 2^{12}$). Results are based on 10^4 independent runs.

Example 1. *Suppose there are* 2 *players,* 2 *arms, and a total of* 2^{16} *rounds. Arm* 1 *is associated with* $\mathcal{B}(0.9)$ *and arm* 2 *is associated with* $\mathcal{B}(0.8)$*. Both players use the* UCB *adaptation described in Section 2. Consider the following three communication strategies:* ***(A)*** *communicate once when* $t = 2^{12}$*;* ***(B)*** *communicate three times when* $t = 2^4, 2^8, 2^{12}$*;* ***(C)*** *communicate* 2^{12} *times when* $t = 1, 2, \ldots, 2^{12}$*.*

Intuitively, Strategy A should be the worst and Strategy C should be the best. However, the numerical experiment gives a surprising result. Figure 1 shows the average number of pulls of the suboptimal arm as a function of time (on a logarithmic scale) for Strategies A, B, C.[2]

To explain this phenomenon, first recall that almost all the single-player regret analysis [1, 2, 5, 6, 10, 13] goes like follows:

> Once a suboptimal arm has been pulled more than $\xi \ln(T)$ times, it will *almost* never be pulled any more before round T.

Here ξ is a constant depending on the bandit algorithm[3]. Following this argument, we first consider Strategy C. It keeps communicating until $t = 2^{12}$. Then according to the full-communication curve, approximately 200 additional explorations of the suboptimal arm are needed before $t = 2^{16}$. However, the two players have to do the 200 explorations **separately** since they cannot communicate any more. Therefore this suboptimal arm is actually explored $2 \cdot 200 = 400$ times from $t = 2^{12}$ to $t = 2^{16}$, resulting in a much higher final regret compared to the full-communication strategy, which only explores the suboptimal arm 200 times during this time period. More generally, after a communication round, if an additional Δ of explorations of the suboptimal arm is needed before the next communication round, the actual explorations performed would be $M \cdot \Delta$. That is, due to lack of communication, $(M-1) \cdot \Delta$ unnecessary explorations are performed, resulting in a larger regret.

On the other hand, Strategy A cleverly makes Δ very close to 0, forcing the final regret almost as low as that of the full-communication curve. When Strategy A finished its only communication, each of the players has already collected approximately 500 independent samples of the suboptimal arm, which is an over-exploration when $t = 2^{12}$ since the full-communication curve indicates that when $t = 2^{12}$, only approximately 300 explorations of the suboptimal arm are needed. However, this amount of exploration happens to be just enough for $t = 2^{16}$ according to the full-communication curve. Therefore, when t goes from 2^{12} to 2^{16}, the suboptimal arm is rarely pulled.

Finally, consider Strategy B. After each communication, the "over-exploration" phenomenon occurs, and the curve acts like that of Strategy A. After some period of time, the amount of explorations

[2] In this experiment, we use $\ln(2t)$ to approximate $\ln(\sum_{k=1}^{K} N_{p,t-1}(k)) + 3\ln\left(\ln\left(\sum_{k=1}^{K} N_{p,t-1}(k)\right)\right)$ for a better comparison when T is relatively small.

[3] e.g., $\frac{1}{2\Delta_a^2}$ in UCB and $\frac{1}{D_{\mathrm{KL}}(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*))}$ in KL-UCB.

goes back to the normal level. Then it will act like Strategy C, resulting in an over-exploration again before the next communication comes. Strategy B performs communication in a relatively stable way. That is, it always keeps Δ relatively small, making its curve fluctuates not so dramatically around the full-communication curve.

Both Strategy A and Strategy B have implied some good ways to develop communication strategies. Strategy A indicates that we can find a time point such that the over-exploration is "just enough." In Section 4 we give the corresponding theoretical guarantees. However, this strategy requires a known time horizon, and the right chance for communication is sort of unpredictable for small T and large M. Fortunately, Strategy B indicates that we can develop anytime strategies that require very few communication rounds with very good performance. The corresponding theoretical guarantees are given in Section 5.1. However, if the communication rounds are too few, there is no way we can use an oblivious bandit policy to achieve a good performance. In Section 5.2, we introduce a non-oblivious bandit policy called DKLUCB (stands for Distributed KL-UCB) which is asymptotically optimal. Most of the results and analyses in Section 4 and Section 5 suppose the KL-UCB adaptation described in Section 2 is used. Similar results also hold for the UCB adaptation (with a different constant) with only slight modification (even simplification) of our analyses. In particular, DKLUCB can be easily modified to get a UCB version called DUCB, which stands for Distributed UCB[4].

4 The Over-Exploration Strategy

In this section, we suppose that a time horizon T is given. That is, the regret is only evaluated at the end of round T. As we have discussed in Section 3, we want to find the right chance for communication such that the resulting over-exploration is "just enough." The following theorem tells us the right chance is around $T^{\frac{1}{M}}$.

Theorem 2. *If the rewards are Bernoulli rewards and the only communication round is round* $\lceil T^{\frac{1}{M}} \rceil$*, then for every suboptimal arm* a*,*

$$\mathbb{E}\left[N_T^{\text{KL-UCB}}(a)\right] \leq \frac{\ln(T)}{D_{\text{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)} + o\left(\ln(T)\right).$$

In other words, using the over-exploration strategy, asymptotically the regret does not scale with the number of players. Strangely, to prove an upper bound, we must prove a lower bound first. In fact, the following lemma is critical to our proof.

Lemma 3. *If the rewards are Bernoulli rewards and the only communication round is round* $\lceil T^{\frac{1}{M}} \rceil$*, then for every suboptimal arm* a *and any* $\delta > 0$*,*

$$\lim_{T\to\infty} \Pr\left(N_{\lceil T^{\frac{1}{M}} \rceil}^{\text{KL-UCB}}(a) \geq \frac{(1-\delta)\ln(T)}{D_{\text{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)}\right) = 1.$$

This lemma ensures that we over-explore the suboptimal arms enough when $t = T^{\frac{1}{M}}$. One way to prove this is to use Theorem 2 in [11]. However, this has some disadvantages. First, it cannot be applied to the UCB adaptation. Second, it requires that we first prove the policy is consistent, which is very unnecessary. Finally, it cannot be translated into a finite-time result. Hence, we present a direct algorithm-oriented proof here. The key idea is, in order to prove something happens with a very low probability, we instead prove that its consequence happens with a very low probability.

Proof sketch of Lemma 3. Let ξ be a shorthand for $1 \,/\, D_{\text{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)$. It suffices to prove that in a single-player setting, $\lim_{T\to\infty} \Pr\left(N_T(a) < (1-\delta)\xi\ln(T)\right) = 0$[5], then use a union bound to get the desired result. Define the random variable Υ_T to be the arm that is pulled most frequently before round T and Ψ to be the last round before round T that Υ_T is pulled. Then

$$\begin{aligned}\Pr\left(N_T(a) < (1-\delta)\xi\ln(T)\right) &\leq \Pr\left(N_T(a) < (1-\delta)\xi\ln(T) \wedge B_{\Psi_T}^+(\Upsilon_T) \geq B_{\Psi_T}^+(a)\right) \\ &\leq \Pr\left(B_{\Psi_T}^+(\Upsilon_T) \geq \mu_{\Upsilon_T} + \epsilon\right) + \Pr\left(\tilde{\mu}_{\Psi_T}(a) \leq \mu_a - \epsilon\right) \quad (4.1) \\ \left(\epsilon \text{ chosen to make this equals } 0\right) + \Pr&\left(D_{\text{KL}}\left(\mathcal{B}(\mu_a - \epsilon)||\mathcal{B}(\mu^* + \epsilon)\right) > \frac{D_{\text{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)}{(1-\delta)}\right).\end{aligned}$$

[4] As explained in [5], UCB is actually a simple relaxation of KL-UCB.

[5] Since we are talking about a single-player setting, here and later, the subscript p is dropped.

By definitions we have $N_{\Psi_T}(\Upsilon_T) \geq (T-1) \,/\, M$, therefore $B^+_{\Psi_T}(\Upsilon_T) \to \tilde{\mu}_{\Psi_T}(\Upsilon_T)$. Now we need to show $N_t(a) \to \infty$ in probability. Then use the fact that $\Psi_T \geq (T-1) \,/\, M$, as well as Hoeffding's inequality, we can prove the two terms in (4.1) converge to 0 as T goes to infinity. □

Then we can start our proof of Theorem 2. The key observation is to define random time points $\Phi_T, \Psi_T, \Lambda_{p,t}$, bound the count after $\Lambda_{p,t}$ using a standard bandit argument, and then bound the count before $\Lambda_{p,t}$ using Lemma 3.

Proof sketch of Theorem 2. Let $\delta > 0$ be an arbitrarily small number, ξ be a shorthand for $1/D_{\mathrm{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)$, and T_0 be a shorthand for $\lceil T^{\frac{1}{M}} \rceil$. We define random variables Φ_T and Ψ_T in the following way: if $N_{T_0}(a) \geq (1-\delta)\xi \ln(T)$, then $\Phi_T = 0$ and $\Psi_T = T_0$; otherwise $\Phi_T = T_0$ and $\Psi_T = T$. We also define random variables $\Lambda_{p,T} = \max\{t : \Phi_T \leq t < \Psi_T, N_{p,t}(a) < (1-\delta)\xi \ln(\Psi_T)\}$. It can be checked that $\Lambda_{p,T}$ is well-defined. By a standard bandit argument, for each player p the expected number of pulls of arm a after $\Lambda_{p,t}$ is no more than $4\delta\xi \ln(T) + o(\ln(T))$, which is negligible. The rest expected number of pulls is no more than

$$\begin{aligned} M \cdot \mathbb{E}\left[N_{1,\Lambda_{1,T}}(a)\right] &\leq M \cdot (\Pr(\Phi_T = 0)(1-\delta)\xi \ln(T_0) + \Pr(\Phi_T = T_0)(1-\delta)\xi \ln(T)) \quad (4.2) \\ \left(\text{by Lemma 3}\right) \quad &\leq \xi(1-\delta)M \ln(\lceil T^{\frac{1}{M}} \rceil) + o(1) \cdot M(1-\delta)\xi \ln(T) \\ &= \xi(1-\delta)\ln(T) + o(\ln(T)). \end{aligned}$$

□

5 Stable Strategies

While the over-exploration strategy works like magic, it has two disadvantages. First, it requires a known time horizon, while in practice we often need an anytime algorithm. Second, when T is relatively small, almost all the bandit policies do much better than the upper bound suggests, which makes the choice of $T^{\frac{1}{M}}$ smaller than actually needed. In other words, Lemma 3 would only make sense when $T^{\frac{1}{M}}$ is relatively large, which may require a huge T if M is large.

To develop anytime algorithms, we first introduce the concept of *communication set*. The communication set is the set of all the communication rounds. Since we want to develop anytime algorithms, we assume the communication set $\mathcal{C}$ is an infinite set whose elements are denoted as $C_1, C_2, C_3, \ldots$ in the increasing order. Then we need to measure the *frequency* of the communication set.

Definition 4. *The counting function on a communication set $\mathcal{C}$ is defined by $\mathcal{Z}_{\mathcal{C}}(n) \coloneqq |\mathcal{C} \cup [n]|$.*

Definition 5. *The density of a communication set $\mathcal{C}$ is defined by $\alpha(\mathcal{C}) \coloneqq \liminf_{k\to\infty} \frac{\ln(C_k)}{\ln(C_{k+1})}$.*

The counting function is a natural and intuitive way to define the frequency of communication. It basically tells us how many communication rounds there are in the first n rounds for any n. However, as we will show later, the *density* is the true property of a communication set in the bandit world. The relationship between these two measurements can be summarized by the following proposition.

Proposition 6. *For every communication set $\mathcal{C}$, **(a)** if $\mathcal{Z}_{\mathcal{C}}(n) \in o(\ln(\ln(n)))$, then $\alpha(\mathcal{C}) = 0$; **(b)** if $\alpha(\mathcal{C}) = 1$, then $\mathcal{Z}_{\mathcal{C}}(n) \in \omega(\ln(\ln(n)))$; **(c)** if $0 < \alpha(\mathcal{C}) < 1$, then*

$$\liminf_{n\to\infty} \frac{\mathcal{Z}_{\mathcal{C}}(n)}{\ln(\ln(n)) \,/\, \ln\left(\alpha(\mathcal{C})^{-1}\right)} \geq 1.$$

In fact, Proposition 6-(a) follows directly from Proposition 6-(b) and (c). As we will show later, a higher density leads to a lower regret. Proposition 6-(b) says that in order to achieve the highest density, or equivalently the lowest regret, the number of communication rounds inevitably falls into the class $\omega(\ln(\ln(n)))$. On the other hand, Proposition 6-(c) says that if we are aiming at a density greater than 0 and less than 1, or equivalently a regret that is "not bad", then the number of communication rounds should be at least in the order of $\ln(\ln(n)) \,/\, \ln(\alpha(\mathcal{C})^{-1})$.

Now we can use the concept of density to establish a lower bound.

Theorem 7. *Let $D_{\inf}(\nu, a, \mathcal{P}) = \inf_{\nu' \in \mathcal{P}: \mathbb{E}[\nu'] > a} D_{\mathrm{KL}}(\nu||\nu')$, $\mathcal{C}$ be the communication set, and π be a consistent bandit policy. Then for every suboptimal arm a satisfying $0 < D_{\inf}(\nu_a, \mu^*, \mathcal{P}) < \infty$,*

$$\limsup_{T\to\infty} \frac{\mathbb{E}[N_T^\pi(a)]}{\ln(T)} \geq \frac{M}{1 + (M-1)\alpha(\mathcal{C})} \cdot \frac{1}{D_{\inf}\left(\nu_a, \mu^*, \mathcal{P}\right)}.$$

To proof Theorem 7, we need to translate any distributed bandit into a single-player bandit, then perform a reduction from the single-player lower bound, and finally calculate out the new constant.

Proof sketch of Theorem 7. First, for each player p, we can use Theorem 1 in [4] to show that

$$\liminf_{t\to\infty} \frac{\mathbb{E}[N_{p,t}(a)]}{\ln(t)} \geq \frac{1}{D_{\inf}(\nu_a, \mu^*, \mathcal{P})}. \tag{5.1}$$

It can be done by relabel the round number in a distributed bandit and translate it into a serialized process, in which a player periodically changes his "role" so that he does not make decision based on all the data available in this serialized process. Then it remains to do some calculation. By the definition of $\alpha(\mathcal{C})$, there exists a subsequence of $(C_k)_{k\geq 1}$, denoted by $(C_{n_s})_{s\geq 1}$, such that

$$\lim_{s\to\infty} \ln(C_{n_s}) \,/\, \ln(C_{n_s+1}) = \alpha(\mathcal{C}). \tag{5.2}$$

Let $x = \limsup_{T\to\infty} \mathbb{E}[N_T(a)]/\ln(T)$. If $x = \infty$, the desired inequality holds trivially. Otherwise,

$$\begin{aligned}
x &\geq \limsup_{s\to\infty} \frac{\mathbb{E}[N_{C_{n_s+1}}(a)]}{\ln(C_{n_s+1})} \geq \liminf_{s\to\infty} \frac{\mathbb{E}\left[\sum_{p=1}^{M} N_{p,C_{n_s+1}-1}(a) - (M-1)\cdot N_{C_{n_s}}(a)\right]}{\ln(C_{n_s+1})} \\
&\geq \sum_{p=1}^{M} \liminf_{s\to\infty} \frac{\mathbb{E}[N_{p,C_{n_s+1}-1}(a)]}{\ln(C_{n_s+1})} - (M-1)\cdot \limsup_{s\to\infty} \frac{\mathbb{E}[N_{C_{n_s}}(a)]}{\ln(C_{n_s+1})} \\
&\geq \sum_{p=1}^{M} \liminf_{t\to\infty} \frac{\mathbb{E}[N_{p,t}(a)]}{\ln(t)} - (M-1)\cdot \limsup_{s\to\infty} \frac{\mathbb{E}[N_{C_{n_s}}(a)]}{\ln(C_{n_s})} \cdot \lim_{s\to\infty} \frac{\ln(C_{n_s})}{\ln(C_{n_s+1})} \\
&\geq \frac{M}{D_{\inf}(\nu_a, \mu^*, \mathcal{P})} - (M-1)\cdot x \cdot \alpha(\mathcal{C}) \qquad \big(\text{by (5.1) and (5.2)}\big) \quad .
\end{aligned}$$

Solving x concludes the proof. □

5.1 Oblivious Policies Under Dense Communication Sets

Theorem 7 shows that if we want to achieve the optimal regret, or in other words, if we do not want the regret scale with the number of players, then the density of the communication set must be 1. Note that linear grid $\{d, 2d, 3d, \dots\}$ and exponential grid $\{q, q^2, q^3, \dots\}$ both have density 1, while a double-exponential grid $\{q^{1+\epsilon}, q^{(1+\epsilon)^2}, q^{(1+\epsilon)^3}, \dots\}$ has density $1/(1+\epsilon) < 1$. The following theorem shows the KL-UCB adaptation achieves this lower bound when $\alpha(\mathcal{C}) = 1$.

Theorem 8. *If the rewards are Bernoulli rewards and the communication set $\mathcal{C}$ satisfies $\alpha(\mathcal{C}) = 1$, then for every suboptimal arm a,*

$$\mathbb{E}\left[N_T^{\text{KL-UCB}}(a)\right] \leq \frac{\ln(T)}{D_{\text{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)} + o\left(\ln(T)\right).$$

Due to the limited space, we defer our proof to Section 5.2, where the DKLUCB policy is introduced. DKLUCB, as a generalization of the KL-UCB adaptation, is optimal even if $\alpha(\mathcal{C}) < 1$.

5.2 Non-Oblivious Policies for Sparse Communication Sets

In Section 5.1 we showed that the KL-UCB adaptation is optimal for dense communication sets (i.e., $\alpha(\mathcal{C}) = 1$). However, if the communication set is very sparse (i.e., $\alpha(\mathcal{C}) < 1$), then we cannot expect an oblivious policy to do uniformly well. As Strategy C in Example 1 has demonstrated, the main difficulty in designing an algorithm for the distributed MAB problem is that each player is "isolated" from others during the period between two communication rounds. In the single-player setting, if one player pulled a suboptimal arm, he should be more certain that this arm is not optimal. Therefore, he should explore this suboptimal arm less frequently in future rounds. However, when it comes to the distributed setting, although each player can utilize the information produced by a suboptimal decision made by himself immediately, other players would not know this experience until next communication round.

Based on the observations above, the key idea is to make each player explore less, since the results of exploration will become common knowledge at the next communication stage. This is implemented by having each player attempt to predict the number of pulls from each arm made by other players since the last communication round. These predictions can be wrong, but not that wrong as all processes are running the same algorithm. In fact, we can show that the errors in these count predictions are negligible as T goes to infinity. However, even if these count predictions were fully correct, we still cannot simulate the full-communication KL-UCB adaptation. This is because the number of the data available is less than the count prediction. For this reason, we have to use a larger confidence bound. The following is a non-oblivious distributed policy called DKLUCB (stands for distributed KL-UCB), where we replace the count $N_{p,t}(a)$ in the KL-UCB adaptation with the count prediction $N'_{p,t}(a)$, and replace the original confidence bound with a slightly larger one.

DKLUCB for Bernoulli rewards. Define $\ell(t) \coloneqq \max\{u \le t : u \in \mathcal{C} \vee u = 0\}$,

$$\mathcal{F}(t) = \frac{M(\ln(t) + 3\ln(\ln(t)))}{1 + (M-1)\alpha(\mathcal{C})},\ u_t(a) = \frac{N_{\ell(t)}(a)}{M}\left(\frac{1}{\alpha(\mathcal{C})} - 1\right),$$

$$N'_{p,t}(a) = N_{p,t}(a) + (M-1)\cdot\min\left(N_{p,t}(a) - N_{\ell(t)}(a), u_t(a)\right).$$

Then choose $A_{p,t} = \arg\max_a B^+_{p,t}(a)$,

$$\text{where } B^+_{p,t}(a) = \sup\left\{p \in (0,1) : D_{\mathrm{KL}}\left(\hat{\nu}_{\mathcal{X}_{p,t-1}(a)} || \mathcal{B}(p)\right) \le \frac{\mathcal{F}\left(\sum_{k=1}^K N_{p,t-1}(k)\right)}{N'_{p,t-1}(a)}\right\}.$$

The DKLUCB policy is a non-oblivious policy because player p in round t makes decision not only based on $\mathcal{X}_{p,t-1}(a)$, but also based on $N_{\ell(t-1)}(a)$, which is the number of pulls of each arm at the end of last communication round. Note that the KL-UCB adaptation can be seen as a special case of the DKLUCB policy. In fact, when $\alpha(\mathcal{C}) = 1$, DKLUCB is identical to the KL-UCB adaptation.

Theorem 9. *If the rewards are Bernoulli rewards and the communication set is $\mathcal{C}$, then for every suboptimal arm a,*

$$\mathbb{E}\left[N_T^{\mathrm{DKLUCB}}(a)\right] \le \frac{M}{1 + (M-1)\alpha(\mathcal{C})} \cdot \frac{\ln(T)}{D_{\mathrm{KL}}\left(\mathcal{B}(\mu_a) || \mathcal{B}(\mu^*)\right)} + o\left(\ln(T)\right).$$

This theorem shows that the DKLUCB policy can achieve the lower bound in Theorem 7. This is consistent with the existing results and intuition. If $M = 1$, then DKLUCB is identical to the single-player KL-UCB policy, and the upper bound is the same. If $\alpha(\mathcal{C}) = 1$, DKLUCB is identical to the KL-UCB adaptation, and the upper bound is still the same (this is formalized as Theorem 8). If $\alpha(\mathcal{C}) = 0$, DKLUCB is no better than M single-player KL-UCB policies running independently.

As in Section 4, to prove an upper bound, first a lower bound is needed. However, this time we need a lemma much stronger than Lemma 3, and theorem 2 in [11] definitely will not help.

Lemma 10. *If the rewards are Bernoulli rewards and the communication set is $\mathcal{C}$, then for every suboptimal arm a and any $\delta > 0$,*

$$\lim_{T\to\infty} \Pr\left(\bigcap_{t \ge T}\left\{N_t^{\mathrm{DKLUCB}}(a) \ge \frac{M}{1 + (M-1)\alpha(\mathcal{C})} \cdot \frac{(1-\delta)\ln(t)}{D_{\mathrm{KL}}(\mathcal{B}(\mu_a) || \mathcal{B}(\mu^*))}\right\}\right) = 1.$$

The idea of the proof is similar to that of Lemma 3, but there are many new ingredients. First we require the following claim to relate the count prediction $N'_{p,t}(a)$ to the true count $N_t(a)$.

Claim 11. *For every a, t, and constant c, if $N_t(a) < c$, then there exists p such that $N'_{p,t}(a) < c$.*

Proof sketch of Lemma 10. Let ξ be a shorthand for the leading constant before $\ln(t)$. It suffices to prove $\lim_{T\to\infty}\Pr(\bigcup_{t\ge T}\{N_t(a) \le (1-\delta)\xi\ln(t)\}) = 0$. For every round t, we define a random player p_t such that $N'_{p_t,t}(a) \le N'_{p',t}(a)$ for every $p' \in [M]$. We say player p_t is *chosen* for round t. We also define random players Γ_t such that Γ_t is the player who is chosen most frequently before round t. Let Ξ_t be a random set including the rounds (before round t) in which player Γ_t is chosen.

Let random variable Υ_t be the arm that is pulled most frequently by Γ_t in those rounds in Ξ_t, and let Ψ_t be the last round in Ξ_t that Γ_t pulls Υ_t. By Claim 11,

$$\Big\{N_t(a) \le (1-\delta)\xi\ln(t)\Big\} \subseteq \Big\{N'_{\Gamma_t,\Psi_t}(a) \le (1-\delta)\xi\ln(t) \wedge B^+_{\Gamma_t,\Psi_t}(\Upsilon_t) \ge B^+_{\Gamma_t,\Psi_t}(a)\Big\}.$$

By definitions of Υ_T, Γ_T, and Ψ_T, we have $\Psi_T \ge (t-1)/(MK)$ and $N_{\Gamma_t,\Psi_t}(\Upsilon_t) \ge (t-1)/(MK)$. The rest is close to the proof of Lemma 3, requiring more careful dealing with the infinite union. □

The following claim explains why we need a slightly larger confidence bound in DKLUCB.

Claim 12. *For every player p, arm a, and $t \ge 0$, $N'_{p,t}(a) \le M \,/\, (1+(M-1)\alpha(\mathcal{C}))\cdot N_{p,t}(a)$.*

In the proof of Theorem 2, choosing $\Lambda_{p,t}$ is relatively easy. However, Theorem 9 requires choosing $\Lambda_{p,T}$ more cleverly. And the overlapping histories cannot be simply ignored like in Theorem 2.

Proof sketch of Theorem 9. Let $\delta > 0$ be an arbitrarily small number and ξ be a shorthand for the leading constant before $\ln(T)$. Define Υ_T to be the largest positive integer such that $C_{\Upsilon_T} \le T$ and $N_{C_{\Upsilon_T}}(a) < (1-\delta)\xi\ln(T)$. Define random variables $\Lambda_{p,T}$ to be the last round such that $C_{\Upsilon_T} \le \Lambda_{p,T} < C_{\Upsilon_T+1}$ and $N'_{p,\Lambda_{p,T}}(a) < (1-\delta)\xi\ln(\min(C_{\Upsilon_T+1}, T))$, or C_{Υ_T} if there is no such round. Using Claim 12, a standard bandit argument would show that for each player p the expected number of pulls of arm a after $\Lambda_{p,t}$ is no more than $4\delta\xi\ln(T)+o(\ln(T))$, which is negligible. Using a decomposition similar to (4.2), by Lemma 10, the rest expected number of pulls is

$$M\cdot\mathbb{E}[N_{1,\Lambda_{1,T}}(a)] - (M-1)\mathbb{E}[N_{\Upsilon_T}(a)] = \mathbb{E}[N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a))]$$
$$\le \mathbb{E}[N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) \,\big|\, N_{C_{\Upsilon_T}}(a) \ge (1-\delta)\xi\ln(T)] + o(\ln(T)).$$

By the definition of $N'_{p,t}(a)$ and $\Lambda_{p,T}$, we have

$$N'_{1,\Lambda_{1,T}}(a) = N_{1,\Lambda_{1,T}}(a) + (M-1)\cdot\min\Big(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a), u_{\Lambda_{1,T}}(a)\Big) < (1-\delta)\xi\ln(T),$$

this plus the definition of $\alpha(\mathcal{C})$ will together imply that, given $N_{C_{\Upsilon_T}}(a) \ge (1-\delta)\xi\ln(T)$,

$$N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) < \xi\ln(T) + o(\ln(T)).$$ □

6 Discussion

Finite-time analysis. We write all the analyses in a way such that they can be translated into finite-time analyses and produce finite-time results (e.g., we avoid using Theorem 2 of [11] to get Lemma 3). However, our results heavily rely on Lemma 3 and Lemma 10, where the technique we used works fine to show the probability converges to 1, yet will produce horrible constant if translated into a finite-time version. Specifically, every time we use pigeonhole principle (e.g., "the arm that is pulled most frequently"), we will add a constant K or M, which we guess is not necessary in the finite-time bound. Our hypothesis is that there are more elegant ways to prove Lemma 3 and Lemma 10, which will provide tighter constants. We hope we can solve this problem in future works.

Difference between bandits and distributed bandits. The typical way to prove an upper bound for traditional bandits is to show that once the upper bound is reached, later suboptimal decisions are negligible. However, this does not work for distributed bandits. The main difficulty is that even if the upper bound is reached *globally* ($N_t(a)$ has reached the upper bound), it may not be reached *locally* ($N_{p,t}(a)$ may be way less than the upper bound). That is why we need the density of communication set, Lemma 3, and Lemma 10 to make sure the desired upper bound is reached locally.

Beyond UCB and KL-UCB. Most of the results and analyses in this paper are very specific, they are either restricted to the KL-UCB adaptation (or the UCB adaptation, after some modification or even simplification, as we have mentioned), or to a generalization of the KL-UCB adaptation (i.e., DKLUCB). However, results similar to Theorem 2, Theorem 8, and Theorem 9 can be reproduced for any UCB-like oblivious bandit policies, as long as they behave normally in the sense that properties similar to Lemma 3 and Lemma 10 hold. These results can also be extended to other oblivious bandit policies that are not based on confidence bound such as Thompson sampling [15]. The intuition behind this is, all the bandit policies behave similarly, and possibly indistinguishable by observing the actions they take. This inspires us to develop a universal framework to prove the performance of bandit policies under distributed settings. However, this framework requires much more insights and we would like to leave it as future work.

References

[1] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.

[2] P. Auer, T. Jaksch, and R. Ortner. Near-optimal regret bounds for reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.

[3] B. Awerbuch and R. Kleinberg. Competitive collaborative learning. *J. Comput. Syst. Sci.*, 74 (8):1271–1288, 2008.

[4] A. Burnetas and M. Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.

[5] O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, and G. Stoltz. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541, 2013.

[6] A. Garivier and O. Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proc. of Conference on Learning Theory (COLT)*, 2011.

[7] E. Hillel, Z. Karnin, T. Koren, R. Lempel, and O. Somekh. Distributed exploration in multi-armed bandits. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.

[8] J. Honda and A. Takemura. An asymptotically optimal bandit algorithm for bounded support models. In *Proc. of Conference on Learning Theory (COLT)*, 2010a.

[9] P. Joulani, A. György, and C. Szepesvári. Online learning under delayed feedback. In *Proc. of International Conference on Machine Learning (ICML)*, 2013.

[10] E. Kaufmann, N. Korda, and R. Munos. Thompson sampling: An asymptotically optimal finite time analysis. *Algorithmic Learning Theory*, 7568:199–213, 2012.

[11] T. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.

[12] K. Liu and Q. Zhao. Distributed learning in multi-armed bandit with multiple players. *IEEE Transactions on Signal Procesing*, 58(11):5667–5681, Nov. 2010.

[13] O.-A. Maillard, R. Munos, and G. Stoltz. A finite-time analysis of multi-armed bandits problems with kullback-leibler divergences. In *Proc. of Conference on Learning Theory (COLT)*, 2011.

[14] B. Szörényi, R. Busa-Fekete, I. Hegedüs, R. Ormándi, M. Jelasity, and B. Kégl. Gossip-based distributed stochastic bandit algorithms. In *Proc. of International Conference on Machine Learning (ICML)*, 2013.

[15] W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.

Appendix A Preliminaries

In this section, we review some notations and introduce new ones that will be used in later proofs. We denote the set of all the positive integers by $\mathbb{N}^+$ and the set of all the positive real numbers by $\mathbb{R}^+$. We denote the set $\{1, 2, 3, \cdots, K\}$ by $[K]$. We denote by $|A|$ the cardinality of a set A. The mean of distribution ν is denoted by $\mathbb{E}[\nu]$. We use $\wedge$ to represent logical conjunction (AND) and use $\vee$ to represent logical disjunction (OR). "$\wedge$" has higher precedence than "$\vee$". Both "$\wedge$" and "$\vee$" have higher precedence than other connectives such as "$=$" or "$\prec$".

A.1 Distributed Bandit Process

The communication set is an infinite set that contains the indices of the communication rounds and is always denoted by $\mathcal{C}$. The elements of $\mathcal{C}$ are denoted by $C_1, C_2, C_3, \ldots$ in the ascending order. We define the function $\ell : \mathbb{N}^+ \to \mathbb{N}$ by

$$\ell(t) \coloneqq \max\{u \leq t : u \in \mathcal{C} \vee u = 0\}.$$

That is, $\ell(t)$ is the last communication round in the first t rounds, and it takes value 0 if there is no such a round. We now define a strict partial order $\prec$ on all the rewards $\mathcal{X} = \{X_{u,v} : u \geq 1, v \geq 1\}$. Concretely, $X_{u_1,v_1} \prec X_{u_2,v_2}$ if and only if

$$v_1 \leq \ell(v_2) \vee u_1 = u_2 \wedge v_1 < v_2.$$

That is, $X_{u_1,v_1} \prec X_{u_2,v_2}$ if and only if anyone who has collected reward X_{u_2,v_2} must also have collected reward X_{u_1,v_1}. For each player p, we define $\prec_p^*$ to be a linear extension of $\prec$ such that $X_{u_1,v_1} \prec_p^* X_{u_2,v_2}$ if and only if

$$v_1 \leq \ell(v_2) \vee \ell(v_1) = \ell(v_2) \wedge (u_1 = p \vee u_1 < u_2)$$

Note the subscript p means this linear extension is defined differently for each player. For player p, this linear extension gives an order on the rewards he receives. It can be checked that both $\prec$ and $\prec_p^*$ are legitimate definitions. We also define

$$\mathcal{X}_{p,t}(a) \coloneqq \{X_{u,v} : X_{u,v} \prec X_{p,t+1} \wedge A_{u,v} = a\} = \{X_{u,v} : X_{u,v} \prec_p^* X_{p,t+1} \wedge A_{u,v} = a\}$$

and $N_{p,t}(a) \coloneqq |\mathcal{X}_{p,t}(a)|$. That is, $\mathcal{X}_{p,t}(a)$ is the set of rewards from arm a that player p has collected at the end of round t, and $N_{p,t}(a)$ is the cardinality of this set.

For each player p and each arm a, we define a sequence of random variables $(X_i^{p,a})_{i\geq 1}$ where $X_i^{p,a}$ is the i_{th} element in the set

$$\{X_{u,v} : X_{u,v} \in \mathcal{X} \wedge A_{u,v} = a\}$$

with respect to the order $\prec_p^*$. We also define

$$\hat{\mu}_{p,s}(a) = \left(\sum_{i=1}^{s} X_i^{p,a}\right) / s.$$

By this definition we can see that $\tilde{\mu}_{p,t}(a) = \hat{\mu}_{p,N_{p,t}(a)}(a)$.

Recall that the random variable $\tilde{\mu}_{p,t}(a)$ is the empirical mean of those rewards generated by arm a that are *known* to player p at the end of round t. There are $N_{p,t}(a)$ of these rewards, each of them obtained either by player p pulling arm a himself, or via communication (i.e., from other players). Thus we need a lemma to ensure that the additional data obtained from other players are indistinguishable from the data collected by players themselves. In other words, we shall prove that $\hat{\mu}_{p,s}(a)$ is the empirical mean of s mutually independent random variables with the same distribution ν_a.

Lemma 13. *For every player p, every arm a, and every $s \in \mathbb{N}^+$,*

$$s \cdot \hat{\mu}_{p,s}(a) \sim B(s, \mu_a),$$

where $B(\cdot,\cdot)$ is a binomial distribution.

Proof. Fix an arm a and a player p. We denote by $F(x)$ the cumulative distribution function corresponding to ν_a. In addition, we define random variables p_i and t_i such that p_i is the player who first receives the reward $X_i^{p,a}$ and t_i is the round this receiving takes place. Clearly we have $X_i^{p,a} = X_{p_i,t_i}$, and therefore

$$s \cdot \hat{\mu}_{p,s}(a) = \sum_{i=1}^{s} X_{p_i,t_i}.$$

Hence, now it suffices to show that X_{p_i,t_i} are mutually independent random variables with the same cumulative distribution function $F(x)$.

First we will prove that for every i, X_{p_i,t_i} has the cumulative distribution function $F(x)$. That is, $\Pr(X_{p_i,t_i} \leq \lambda) = F(\lambda)$ for any $\lambda \in \mathbb{R}$. Note that

$$\begin{aligned}
\Pr(X_{p_i,t_i} \leq \lambda) &= \sum_{u,v} \Pr\left(X_{u,v} \leq \lambda \wedge p_i = u \wedge t_i = v\right) \\
&= \sum_{u,v:\Pr(p_i=u\wedge t_i=v)>0} \Pr\left(X_{u,v} \leq \lambda \mid p_i = u \wedge t_i = v\right) \Pr\left(p_i = u \wedge t_i = v\right) \\
&= \sum_{u,v:\Pr(p_i=u\wedge t_i=v)>0} F(\lambda) \Pr\left(p_i = u \wedge t_i = v\right) = F(\lambda) \qquad \text{(A.1)}
\end{aligned}$$

where the second equality holds because given $p_i = u$ and $t_i = v$, the reward $X_{u,v}$ is generated by arm a independently.

Then we will prove mutual independence by induction. It suffices to show that for every s real numbers $\lambda_1, \lambda_2, \cdots, \lambda_s$ and s integers $l_1 < l_2 < \cdots < l_s$ we have

$$\Pr\left(\bigwedge_{i=1}^{s}(X_{p_{l_i},t_{l_i}} \leq \lambda_i)\right) = \prod_{i=1}^{s} F(\lambda_i).$$

The base case where $s = 1$ has been proved in (A.1). Suppose we have proved that mutual independence holds for $s-1$, and we also define the event $\mathcal{E}_{s,u,v}$ by

$$\mathcal{E}_{s,u,v} \coloneqq \left\{\bigwedge_{i=1}^{s-1}(X_{p_{l_i},t_{l_i}} \leq \lambda_i) \wedge p_{l_s} = u \wedge t_{l_s} = v\right\}.$$

Then we have

$$\begin{aligned}
\Pr\left(\bigwedge_{i=1}^{s}(X_{p_{l_i},t_{l_i}} \leq \lambda_i)\right) &= \sum_{u,v} \Pr\left(X_{u,v} \leq \lambda_s \wedge \mathcal{E}_{s,u,v}\right) \\
&= \sum_{u,v:\Pr(\mathcal{E}_{s,u,v})>0} \Pr\left(X_{u,v} \leq \lambda_s \mid \mathcal{E}_{s,u,v}\right) \Pr\left(\mathcal{E}_{s,u,v}\right) \\
&= \sum_{u,v:\Pr(\mathcal{E}_{s,u,v})>0} F(\lambda_s) \Pr\left(\mathcal{E}_{s,u,v}\right) \\
&= F(\lambda_s) \prod_{i=1}^{s-1} \Pr\left(\bigwedge_{i=1}^{s-1}(X_{p_{l_i},t_{l_i}} \leq \lambda_i)\right) \\
&= F(\lambda_s) \prod_{i=1}^{s-1} F(\lambda_i) \qquad \text{(by induction hypothesis)} \\
&= \prod_{i=1}^{s} F(\lambda_i),
\end{aligned}$$

where the third equality holds because given the event $\mathcal{E}_{s,u,v}$, the reward $X_{u,v}$ is generated by arm a independently. □

Technically, this lemma is required whenever the Hoeffding's inequality is used to bound $\hat{\mu}_{p,s}(a)$. For simplicity, later proofs may use this lemma **without explicitly pointing it out**.

A.2 Kullback-Leibler Divergences

The Kullback-Leibler divergence (KL-divergence) from probability distribution ν_1 to probability distribution ν_2 is defined by

$$D_{\mathrm{KL}}(\nu_1||\nu_2) \coloneqq -\mathbb{E}_{\nu_1}\left[\ln\left(\frac{d\nu_2}{d\nu_1}\right)\right].$$

Accordingly, we define

$$D_{\inf}(\nu, a, \mathcal{P}) \coloneqq \inf_{\nu' \in \mathcal{P}:\mathbb{E}[\nu']>a} D_{\mathrm{KL}}(\nu||\nu').$$

Clearly, for two Bernoulli distribution ν_1 and ν_2 satisfying $\mathbb{E}[\nu_1] < \mathbb{E}[\nu_2]$, we have

$$D_{\mathrm{KL}}(\nu_1||\nu_2) = \mathcal{K}_{\inf}(\nu_1, \mathbb{E}[\nu_2], \mathcal{B}),$$

where $\mathcal{B}$ is the set of all the Bernoulli distributions. Note that the parameter of a Bernoulli distribution usually takes value in the open interval $(0, 1)$. However, the empirical mean of Bernoulli trials can take value in the closed interval $[0, 1]$. Hence we define the extended Bernoulli distribution with parameter $p \in [0, 1]$ to be a distribution having probability mass p on 1 and $1 - p$ on 0. We let

$$\mathcal{K}(p, q) \coloneqq p \ln\left(\frac{p}{q}\right) + (1-p)\ln\left(\frac{1-p}{1-q}\right)$$

be the KL-divergence from an extended Bernoulli distribution with parameter p to another with parameter q, with conventions $0 \cdot \ln(0) = 0$ and $\ln(0\,/\,0) = 0$. We also define the left-side truncated KL-divergence $\mathcal{K}'(p, q)$ as 0 if $p > q$, or $\mathcal{K}(p, q)$ otherwise.

A.3 Tools From Single-Player Bandits

The original proof of the KL-UCB algorithm mainly relies on the following self normalized deviation bound, which we cannot avoid either.

Lemma 14 ([5])**.** *Let $\hat{\mu}_s$ be the empirical mean of s mutually independent Bernoulli random variables with the same mean p, then*

$$\Pr\left(\bigcup_{s=1}^{t}\left\{\hat{\mu}_s < p \wedge s\cdot\mathcal{K}\left(\hat{\mu}_s, p\right) \geq \epsilon\right\}\right) \leq e\lceil\epsilon\ln(t)\rceil e^{-\epsilon}.$$

Appendix B Proof of Theorem 2

Let $\delta > 0$ be an arbitrarily small number, ξ be a shorthand for $1/D_{\mathrm{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)$, and T_0 be a shorthand for $\lceil T^{\frac{1}{M}}\rceil$. We define random variables Φ_T, Ψ_T in the following way: if $N_{T_0}(a) \geq (1-\delta)\xi\ln(T)$, then $\Phi_T = 0$ and $\Psi_T = T_0$; otherwise $\Phi_T = T_0$ and $\Psi_T = T$. We also define random variables $\Lambda_{p,T} = \max\{t : \Phi_T \leq t < \Psi_T, N_{p,t}(a) < (1-\delta)\xi\ln(\Psi_T)\}$. It can be checked that $\Lambda_{p,T}$ is well-defined.

Step 1: Bound the count after $\Lambda_{p,T}$ (traditional bandit argument). We first do an event decomposition:

$$\{A_{p,t} = a\} \subseteq \left\{B^+_{p,t}(a^*) < \mu^*\right\} \cup \left\{B^+_{p,t}(a) \geq \mu^* \wedge A_{p,t} = a\right\}, \quad \text{for } t \text{ large enough.}$$

Then we will show that the event $\left\{B^+_{p,t}(a^*) < \mu^*\right\}$ can be safely ignored.

Lemma 15. *For every player p and every arm a,*

$$\sum_{t=1}^{T}\Pr\left(B^+_{p,t}(a) < \mu_a\right) = o\left(\ln(T)\right).$$

Proof. Note that

$$\begin{aligned}
\Pr\left(B^+_{p,t}(a) < \mu_a\right) &\leq \Pr\left(N_{p,t}(a)\mathcal{K}'(\hat{\mu}_{p,N_{p,t}(a)}(a), \mu_a) \geq \ln(t) + 3\ln(\ln(t))\right) \\
&\leq \Pr\left(\bigcup_{s=1}^{Mt}\left\{\hat{\mu}_{p,s}(a) < \mu_a \wedge s\mathcal{K}(\hat{\mu}_{p,s}(a), \mu_a) \geq \ln(t) + 3\ln(\ln(t))\right\}\right) \\
&\leq \frac{e\lceil(\ln(t) + 3\ln(\ln(t)))\ln(Mt)\rceil}{t\ln^3(t)}, \qquad \text{(B.1)}
\end{aligned}$$

where the last inequality follows from Lemma 14. Sum (B.1) from 1 to T yields $o(\ln(T))$. □

Hence we can ignore the event $\left\{B^+_{p,t}(a^*) < \mu^*\right\}$ and only bound the probability of the event $\left\{B^+_{p,t}(a) \geq \mu^* \wedge A_{p,t} = a\right\}$.

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=\Lambda_T+1}^{T}\mathbb{1}_{\left\{B^+_{p,t}(a)\geq\mu^*\wedge A_{p,t}=a\right\}}\right] \\
&\leq \mathbb{E}\left[\sum_{t=\Lambda_T+1}^{T}\mathbb{1}_{\left\{\mathcal{K}'(\tilde{\mu}_{p,t}(a),\mu^*)\leq\frac{\mathcal{K}(\mu_a,\mu^*)}{1+\delta}\wedge A_{p,t}=a\right\}}\right] + \mathbb{E}\left[\sum_{t=\Lambda_T+1}^{T}\mathbb{1}_{\{N_{p,t}(a)<(1+\delta)\xi\ln(t)\wedge A_{p,t}=a\}}\right] + o(\ln(T)) \\
&\leq \mathbb{E}\left[\sum_{s=0}^{\infty}\mathbb{1}_{\left\{\mathcal{K}'(\hat{\mu}_{p,s}(a),\mu^*)\leq\frac{\mathcal{K}(\mu_a,\mu^*)}{1+\delta}\right\}}\right] + 4\delta\xi\ln(T) + o(\ln(T)) \quad \text{(by the definition of } \Lambda_{p,T}\text{)} \\
&\leq \sum_{s=0}^{\infty}\Pr\left(\hat{\mu}_{p,s}(a) \geq \mu + \epsilon\right) + 4\delta\xi\ln(T) + o(\ln(T)) \quad \text{(for some } \epsilon > 0\text{)} \\
&\leq \sum_{s=0}^{\infty}e^{-2\epsilon^2 s} + 4\delta\xi\ln(T) + o(\ln(T)) \quad \text{(by Hoeffding's inequality)} \\
&= 4\delta\xi\ln(T) + o(\ln(T)),
\end{aligned}$$

Therefore, for all the players, the count after $\Lambda_{p,T}$ is no more than $4M\delta\xi\ln(T) + o(\ln(T))$.

Step 2: Bound the count before $\Lambda_{p,T}$. The total count of all players before $\Lambda_{p,T}$ should be no more than

$$\begin{aligned}
&\sum_{p=1}^{M} \mathbb{E}\left[N_{p,\Lambda_{p,T}}(a)\right] - (M-1)\mathbb{E}\left[N_{\Phi_T}(a)\right] \\
&\leq M \cdot \mathbb{E}\left[N_{1,\Lambda_{1,T}}(a)\right] \\
&\leq M \cdot (\Pr(\Phi_T = 0)(1-\delta)\xi\ln(T_0) + \Pr(\Phi_T = T_0)(1-\delta)\xi\ln(T)) \\
\text{(by Lemma 3)} \quad &\leq \xi(1-\delta)M\ln(\lceil T^{\frac{1}{M}}\rceil) + o(1)\cdot M(1-\delta)\xi\ln(T) \\
&= \xi(1-\delta)\ln(T) + o(\ln(T)).
\end{aligned}$$

Step 3: Put everything together. Adding up all the components, we get

$$\mathbb{E}\left[N_T(a)\right] \leq \xi\left(1 + (4M-1)\delta\right)\ln(T) + o(\ln(T)).$$

This concludes the proof.

Appendix C Proof of Lemma 3

Let ξ be a shorthand for $1 \,/\, D_{\mathrm{KL}}\left(\mathcal{B}(\mu_a)||\mathcal{B}(\mu^*)\right)$. Then it suffices to prove that in a single player setting, $\lim_{T\to\infty}\Pr\left(N_T(a) \leq (1-\delta)\xi\ln(T)\right) = 1$. and then use a union bound.

We can see this as weaker form of a special case of Lemma 10. It is weaker because it does not contain a infinite intersection like Lemma 10. It is a special case becuase we can let $M = 1$ in Lemma 10. For these reasons, the proof should be a simplified version of the proof of Lemma 10. To avoid duplication, we refer the reader to the proof in Section G.

Appendix D Proof of Proposition 6

Since Proposition 6-(a) is a direct consequence of Proposition 6-(b) and (c), it suffices to prove the latter two statements. Let

$$\tilde{\mathcal{C}} = \{\ln(C_1), \ln(C_2), \ln(C_3), \cdots\}$$

and we denote $\ln(C_k)$ by $\tilde{C}_k$. Assume $\alpha(\mathcal{C}) = d \in (0,1]$. Then for every $\epsilon \in (0,d)$, there exists a N_ϵ large enough such that $\tilde{C}_k \,/\, \tilde{C}_{k+1} \geq d - \epsilon$ for every $k \geq N_\epsilon$. Thus, for n large enough,

$$\ln(n) < \tilde{C}_{\mathcal{Z}_{\mathcal{C}}(n)+1} \leq \frac{\tilde{C}_{N_\epsilon}}{(d-\epsilon)^{\mathcal{Z}_{\mathcal{C}}(n)+1-N_\epsilon}}.$$

It then follows from the inequality that

$$\mathcal{Z}_{\mathcal{C}}(n) > \frac{\ln(\ln(n)) - \ln(\tilde{C}_{N_\epsilon})}{\ln(\frac{1}{d-\epsilon})} + N_\epsilon - 1.$$

Note that ϵ can be chosen to be arbitrarily small. If $d = 1$, then $\mathcal{Z}_{\mathcal{C}}(n) \in \omega(\ln\ln(n))$. Otherwise, we have

$$\liminf_{n\to\infty} \frac{\mathcal{Z}_{\mathcal{C}}(n)}{\ln(\ln(n)) \,/\, \ln\left(d^{-1}\right)} \geq 1.$$

This completes the proof of Proposition 6-(b) and (c).

Appendix E Proof of Theorem 7

For the single-player MAB model, [11] gave a lower bound for single-parametric distributions. [4] generalized this result to non-parametric models. Translated into our model, their results can be summarized as the following theorem.

Theorem 16 ([4]). *If $M = 1$ and π is a consistent policy, then for every suboptimal arm a satisfying $0 < D_{\inf}(\nu_a, \mu^*, \mathcal{P}) < \infty$,*

$$\liminf_{T\to\infty} \frac{\mathbb{E}[N_T^\pi(a)]}{\ln(T)} \geq \frac{1}{D_{\inf}(\nu_a, \mu^*, \mathcal{P})}.$$

Now our goal is to establish a similar lower bound for the distributed case (i.e., $M > 1$). Fortunately, we can actually use Theorem 16 as a stepping stone. More specifically, we could use a *simulator* to simulate all the actions of the M players and apply Theorem 16 to this single simulator. In other words, we treat every multi-player MAB process as a single-player MAB process whose outcome is indistinguishable from the original one. This conversion does not provide a direct solution to the lower bound for the distributed MAB problem. However, we can use it to establish a lower bound on $N_{p,t}(a)$ for every suboptimal arm a.

Lemma 17. *If π is a consistent policy, then for every suboptimal arm a satisfying $0 < D_{\inf}(\nu_a, \mu^*, \mathcal{P}) < \infty$,*

$$\liminf_{t\to\infty} \frac{\mathbb{E}[N_{p,t}^\pi(a)]}{\ln(t)} \geq \frac{1}{D_{\inf}(\nu_a, \mu^*, \mathcal{P})}.$$

Proof. Fix a player p and a suboptimal arm a. Let $\mathcal{D}$ be the original distributed process. The rewards in the set $\mathcal{X}$ can be viewed as ones generated by a single-player MAB process $\mathcal{S}$ in order $\prec_p^*$ using a single-player policy π'. Now we use coupling to associate process $\mathcal{S}$ to process $\mathcal{D}$ and use superscripts to distinguish random variables in the two processes. Note that

$$N_t^{\mathcal{S}}(a) \leq N_{p,t}^{\mathcal{D}}(a) \leq N_t^{\mathcal{D}}(a). \tag{E.1}$$

Note that the subscript t in $N_t^{\mathcal{S}}(a)$ is referring to the time point t in the serialized process, while the subscript t in $N_{p,t}^{\mathcal{D}}(a)$ and $N_t^{\mathcal{D}}(a)$ are referring to the time point t in the distributed process.

Hence, if π is a consistent policy (for the distributed MAB), then π' is a consistent policy (for the single-player MAB). Thus,

$$\liminf_{t\to\infty} \frac{\mathbb{E}[N_{p,t}^{\mathcal{D}}(a)]}{\ln(t)} \geq \liminf_{t\to\infty} \frac{\mathbb{E}[N_t^{\mathcal{S}}(a)]}{\ln(t)} \geq \frac{1}{D_{\inf}(\nu_a, \mu^*, \mathcal{P})}.$$

where the first inequality follows from (E.1) and the second follows from Theorem 16. □

Having a lower bound on $N_{p,t}(a)$ is almost equivalent to having a lower bound on $N_t(a)$. In fact, we have for every $t \geq 1$,

$$N_t(a) \geq N_{\ell(t)}(a) + \sum_{p=1}^{M} \big(N_{p,t}(a) - N_{\ell(t)}(a)\big) = \sum_{p=1}^{M} N_{p,t}(a) - (M-1)\cdot N_{\ell(t)}(a). \tag{E.2}$$

Using (E.2) we can finish the proof of Theorem 7. Since $\alpha(\mathcal{C}) = \liminf_{k\to\infty} \frac{\ln(C_k)}{\ln(C_{k+1})}$, there exists a subsequence of $(C_k)_{k\geq 1}$, denoted by $(C_{n_s})_{s\geq 1}$, such that

$$\lim_{s\to\infty} \frac{\ln(C_{n_s})}{\ln(C_{n_s+1})} = \alpha(\mathcal{C}).$$

Let $x = \limsup_{T\to\infty} \frac{\mathbb{E}[N_T(a)]}{\ln(T)} \geq 0$. If $x = \infty$, the desired inequality holds trivially. Assume x is finite. Then

$$\begin{aligned}
x &\geq \limsup_{s\to\infty} \frac{\mathbb{E}[N_{C_{n_s+1}}(a)]}{\ln(C_{n_s+1})} \geq \liminf_{s\to\infty} \frac{\mathbb{E}\left[\sum_{p=1}^{M} N_{p,C_{n_s+1}-1}(a) - (M-1)\cdot N_{C_{n_s}}(a)\right]}{\ln(C_{n_s+1})} \\
&\geq \sum_{p=1}^{M} \liminf_{s\to\infty} \frac{\mathbb{E}[N_{p,C_{n_s+1}-1}(a)]}{\ln(C_{n_s+1})} - (M-1)\cdot \limsup_{s\to\infty} \frac{\mathbb{E}[N_{C_{n_s}}(a)]}{\ln(C_{n_s+1})} \\
&\geq \sum_{p=1}^{M} \liminf_{t\to\infty} \frac{\mathbb{E}[N_{p,t-1}(a)]}{\ln(t)} - (M-1)\cdot \limsup_{s\to\infty} \frac{\mathbb{E}[N_{C_{n_s}}(a)]}{\ln(C_{n_s})} \cdot \lim_{s\to\infty} \frac{\ln(C_{n_s})}{\ln(C_{n_s+1})} \\
&\geq \frac{M}{D_{\inf}(\nu_a, \mu^*, \mathcal{P})} - (M-1)\cdot x \cdot \alpha(\mathcal{C}). \qquad \text{(by Lemma 17)}
\end{aligned}$$

Solving x concludes the proof.

Appendix F Proof of Theorem 8 and Theorem 9

Theorem 8 is a special case of Theorem 9: If $\alpha(\mathcal{C}) = 1$, then the DKLUCB algorithm reduces to a simple adaptation of the KL-UCB, and the upper bound is identical to the single-player upper bound. Therefore, it suffices to prove Theorem 9.

Let $\delta > 0$ be an arbitrarily small number and ξ be a shorthand for the leading constant before $\ln(T)$. Define Υ_T to be the largest positive integer such that $C_{\Upsilon_T} \le T$ and $N_{C_{\Upsilon_T}}(a) < (1-\delta)\xi \ln(T)$. Define random variables $\Lambda_{p,T}$ to be the last round such that $C_{\Upsilon_T} \le \Lambda_{p,T} < C_{\Upsilon_T+1}$ and $N'_{p,\Lambda_{p,T}}(a) < (1-\delta)\xi \ln(\min(C_{\Upsilon_T+1}, T))$. If there is no such round, let $\Lambda_{p,T} = C_{\Upsilon_T}$.

Step 1: bound the count after $\Lambda_{p,T}$ (traditional bandit argument).

We first do an event decomposition:

$$\{A_{p,t} = a\} \subseteq \left\{B^+_{p,t}(a^*) < \mu^*\right\} \cup \left\{B^+_{p,t}(a) \ge \mu^* \wedge A_{p,t} = a\right\}, \quad \text{for } t \text{ large enough.}$$

Then we will show that the event $\left\{B^+_{p,t}(a^*) < \mu^*\right\}$ can be safely ignored.

Lemma 18. *For every player p and every arm a,*

$$\sum_{t=1}^{T} \Pr\left(B^+_{p,t}(a) < \mu_a\right) = o\left(\ln(T)\right).$$

Proof. Note that

$$\begin{aligned}
\Pr\left(B^+_{p,t}(a) < \mu_a\right) &\le \Pr\left(N'_{p,t}(a)\mathcal{K}'(\hat{\mu}_{p,N_{p,t}(a)}(a), \mu_a) \ge \xi\left(\ln(t) + 3\ln(\ln(t))\right)\right) \\
\left(\text{by Claim 12}\right) &\le \Pr\left(N_{p,t}(a)\mathcal{K}'(\hat{\mu}_{p,N_{p,t}(a)}(a), \mu_a) \ge \ln(t) + 3\ln(\ln(t))\right) \\
&\le \Pr\left(\bigcup_{s=1}^{Mt} \{\hat{\mu}_{p,s}(a) < \mu_a \wedge s\mathcal{K}(\hat{\mu}_{p,s}(a), \mu_a) \ge \ln(t) + 3\ln(\ln(t))\}\right) \\
&\le \frac{e\lceil(\ln(t) + 3\ln(\ln(t)))\ln(Mt)\rceil}{t\ln^3(t)}, \qquad \text{(F.1)}
\end{aligned}$$

where the last inequality follows from Lemma 14. Sum (F.1) from 1 to T yields $o(\ln(T))$. □

Hence we can ignore the event $\left\{B^+_{p,t}(a^*) < \mu^*\right\}$ and only bound the probability of the event $\left\{B^+_{p,t}(a) \ge \mu^* \wedge A_{p,t} = a\right\}$.

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=\Lambda_T+1}^{T} \mathbb{1}_{\left\{B^+_{p,t}(a) \ge \mu^* \wedge A_{p,t} = a\right\}}\right] \\
&\le \mathbb{E}\left[\sum_{t=\Lambda_T+1}^{T} \mathbb{1}_{\left\{\mathcal{K}'(\tilde{\mu}_{p,t}(a), \mu^*) \le \frac{\mathcal{K}(\mu_a, \mu^*)}{1+\delta} \wedge A_{p,t} = a\right\}}\right] + \mathbb{E}\left[\sum_{t=\Lambda_T+1}^{T} \mathbb{1}_{\left\{N'_{p,t}(a) < (1+\delta)\xi\ln(t) \wedge A_{p,t} = a\right\}}\right] + o(\ln(T)) \\
&\le \mathbb{E}\left[\sum_{s=0}^{\infty} \mathbb{1}_{\left\{\mathcal{K}'(\hat{\mu}_{p,s}(a), \mu^*) \le \frac{\mathcal{K}(\mu_a, \mu^*)}{1+\delta}\right\}}\right] + 4\delta\xi\ln(T) + o(\ln(T)) \quad \text{(by the definition of } \Lambda_{p,T}\text{)} \\
&\le \sum_{s=0}^{\infty} \Pr\left(\hat{\mu}_{p,s}(a) \ge \mu + \epsilon\right) + 4\delta\xi\ln(T) + o(\ln(T)) \quad \text{(for some } \epsilon > 0\text{)} \\
&\le \sum_{s=0}^{\infty} e^{-2\epsilon^2 s} + 4\delta\xi\ln(T) + o(\ln(T)) \quad \text{(by Hoeffding's inequality)} \\
&= 4\delta\xi\ln(T) + o(\ln(T)),
\end{aligned}$$

Therefore, for all the players, the count after $\Lambda_{p,T}$ is no more than $4M\delta\xi\ln(T) + o(\ln(T))$.

Step 2: bound the count before $\Lambda_{p,T}$. The total count of all players before $\Lambda_{p,T}$ should be no more than

$$\begin{aligned}
&\sum_{p=1}^{M} \mathbb{E}\left[N_{p,\Lambda_{p,T}}(a)\right] - (M-1)\mathbb{E}\left[N_{C_{\Upsilon_T}}(a)\right] \\
&= M \cdot \mathbb{E}\left[N_{1,\Lambda_{1,T}}(a)\right] - (M-1)\mathbb{E}\left[N_{C_{\Upsilon_T}}(a)\right] \\
&= \mathbb{E}\left[N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a))\right] \\
&\le \Pr\left(N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(C_{\Upsilon_T})\right) \cdot \\
&\qquad \mathbb{E}\left[N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) \mid N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(C_{\Upsilon_T})\right] \\
&\quad + \Pr\left(N_{C_{\Upsilon_T}} < (1-\delta)\xi\ln(C_{\Upsilon_T})\right) \cdot \\
&\qquad \mathbb{E}\left[N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) \mid N_{C_{\Upsilon_T}} < (1-\delta)\xi\ln(C_{\Upsilon_T})\right] \\
&\le \mathbb{E}\left[N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) \mid N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(C_{\Upsilon_T})\right] \qquad \text{(F.2)} \\
&\qquad + o(\ln(T)) \qquad \text{(by Lemma 10)}
\end{aligned}$$

Our goal is to prove (F.2) is no more than $\xi\ln(T)+o(\ln(T))$. It can be done by showing $N_{1,\Lambda_{1,T}}(a)+(M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) \le \xi\ln(T) + o(\ln(T))$ given $N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(C_{\Upsilon_T})$. Now we suppose $N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(C_{\Upsilon_T})$ holds. If $N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}} \le u_{\Lambda_{1,T}}(a)$, then it is trivial since in that case $N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) = N'_{1,\Lambda_{1,T}}(a)$ and we already know $N'_{1,\Lambda_{1,T}}(a) < (1-\delta)\xi\ln(T)$. Otherwise we have

$$N_{1,\Lambda_{1,T}}(a) + (M-1)\frac{N_{C_{\Upsilon_T}}(a)}{M}\left(\frac{1}{\alpha(\mathcal{C})} - 1\right) < (1-\delta)\xi\ln\left(\min(T, C_{\Upsilon_T+1})\right).$$

Using the condition $N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(C_{\Upsilon_T})$ we get

$$N_{1,\Lambda_{1,T}}(a) < (1-\delta)\xi\left(\ln\left(\min(T, C_{\Upsilon_T+1})\right) - \frac{M-1}{M}\left(\frac{\ln(C_{\Upsilon_T})}{\alpha(\mathcal{C})} - \ln(C_{\Upsilon_T})\right)\right).$$

By the definition of $\alpha(\mathcal{C})$ we have

$$\begin{aligned}
N_{1,\Lambda_{1,T}}(a) &< (1-\delta)\xi\left((1+\delta)\frac{\ln\left(\min(T, C_{\Upsilon_T+1})\right)}{M} + \frac{M-1}{M}\ln(C_{\Upsilon_T})\right) + o(\ln(T)) \\
&< \frac{\xi\ln(T)}{M} + \frac{(1-\delta)\xi(M-1)}{M}\ln(C_{\Upsilon_T}) + o(\ln(T)). \qquad \text{(F.3)}
\end{aligned}$$

Therefore, given $N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(C_{\Upsilon_T})$,

$$\begin{aligned}
&N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) \\
&= M \cdot N_{1,\Lambda_{1,T}}(a) - (M-1)\cdot N_{C_{\Upsilon_T}}(a) \\
&\le M \cdot N_{1,\Lambda_{1,T}}(a) - (M-1)(1-\delta)\xi\ln(C_{\Upsilon_T}) \\
&\le \xi\ln(T) + o(\ln(T)) \qquad \text{(by (F.3))}
\end{aligned}$$

Hence,

$$\mathbb{E}\left[N_{1,\Lambda_{1,T}}(a) + (M-1)(N_{1,\Lambda_{1,T}}(a) - N_{C_{\Upsilon_T}}(a)) \mid N_{C_{\Upsilon_T}} \ge (1-\delta)\xi\ln(T)\right] \le \xi\ln(T) + o(\ln(T))$$

Step 3: put everything together. Adding all components up, we get

$$\mathbb{E}[N_T(a)] \le \xi(1+4M\delta)\ln(T) + o(\ln(T)).$$

This concludes the proof.

Appendix G Proof of Lemma 10

In this section, in order to make the proof easier and more readable, we will prove a lemma equivalent to Lemma 10, We first introduce some new concepts.

Definition 19. *We say a sequence of random variables* $(X_n)_{n\geq 1}$

(a) converges to a constant c *in probability, denoted by* $X_n \xrightarrow{p} c$*, if for every* $\epsilon > 0$*,*

$$\lim_{n\to\infty} \Pr\left(|X_n - c| \geq \epsilon\right) = 0;$$

(b) tends to infinity in probability, denoted by $X_n \xrightarrow{p} \infty$*, if for every* N*,*

$$\lim_{n\to\infty} \Pr\left(X_n < N\right) = 0.$$

The following is the equivalent lemma we will prove.

Lemma 20 (equivalent to Lemma 10)**.** *Let* $(\Phi_n)_{n\geq 1}$ *be a sequence of random variables such that* $\Phi_n \xrightarrow{p} \infty$*. Then for every suboptimal arm* a *and every* $\delta > 0$*,*

$$\lim_{n\to\infty} \Pr\left(N_{\Phi_n}(a) \geq \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})}\right) = 1.$$

Note that this is a classical technique to deal with infinite union (or intersection), we omit the proof of equivalence here.

To simplify our proof, we first present three utility lemmas.

Lemma 21. *Let* $(\Phi_n)_{n\geq 1}$ *be a sequence of random variables such that* $\Phi_n \xrightarrow{p} \infty$*. Then for every player* p *and every arm* a*,*

$$\hat{\mu}_{p,\Phi_n}(a) \xrightarrow{p} \mu_a.$$

Proof. By the definition of convergence in probability, it suffices to show that for every $\delta > 0$ and every $\epsilon > 0$ we can find an N_* such that for any $n \geq N_*$,

$$\Pr\left(|\hat{\mu}_{p,\Phi_n}(a) - \mu_a| \geq \delta\right) \leq \epsilon.$$

Fix a $\delta > 0$ and $\epsilon > 0$, we can choose N_0 large enough such that

$$\Pr\left(\bigcup_{s=N_0}^{\infty} |\hat{\mu}_{p,s}(a) - \mu_a| \geq \delta\right) \leq \sum_{s=N_0}^{\infty} \Pr\left(|\hat{\mu}_{p,s}(a) - \mu_a| \geq \delta\right) \leq \sum_{s=N_0}^{\infty} 2e^{-2\delta^2 s} \leq \frac{\epsilon}{2}.$$

where the first inequality follows from union bound and the second follows from Hoeffding's inequality. Then by the definition of tending to infinity in probability, we can choose N_1 large enough such that

$$\Pr\left(\Phi_n < N_0\right) \leq \frac{\epsilon}{2}, \text{ for every } n \geq N_1.$$

Thus, for every $n \geq N_* = \max(N_0, N_1)$,

$$\Pr\left(|\hat{\mu}_{p,\Phi_n}(a) - \mu_a| \geq \delta\right) \leq \Pr\left(\Phi_n < N_0\right) + \Pr\left(\bigcup_{s=N_0}^{\infty} |\hat{\mu}_{p,s}(a) - \mu_a| \geq \delta\right) \leq \epsilon,$$

which concludes the proof. □

Lemma 22. *Let* $(\Upsilon_n)_{n\geq 1}$ *be a sequence of random arms,* $(\Gamma_n)_{n\geq 1}$ *be a sequence of random players, and* $(\Phi_n)_{n\geq 1}$ *be a sequence of random variables such that* $\Phi_n \xrightarrow{p} \infty$*. Then*

$$\exists \xi > 0, N_0 > 0, \forall n \geq N_0, N_{\Gamma_n,\Phi_n}(\Upsilon_n) \geq (\Phi_n)^{\xi} \implies \left(B^{+}_{\Gamma_n,\Phi_n}(\Upsilon_n) - \mu_{\Upsilon_n}\right) \xrightarrow{p} 0. \quad \text{(G.1)}$$

Proof. Assume the left hand side of (G.1) holds. By the definition of convergence in probability, it suffices to show that for every $\delta > 0$ we have

$$\lim_{n\to\infty} \Pr\left(|B^+_{\Gamma_n,\Phi_n}(\Upsilon_n) - \mu_{\Upsilon_n}| \geq \delta\right) = 0.$$

Fix a $0 < \delta < 1 - \max_{a\in\mathcal{A}} \mu_a$, note that

$$\begin{aligned}
\Pr\left(|B^+_{\Gamma_n,\Phi_n}(\Upsilon_n) - \mu_{\Upsilon_n}| > \delta\right) &\leq \Pr\left(\hat{\mu}_{\Gamma_n,N_{\Gamma_n,\Phi_n}(\Upsilon_n)}(\Upsilon_n) < \mu_{\Upsilon_n} - \delta\right) \\
&+ \Pr\left(\hat{\mu}_{\Gamma_n,N_{\Gamma_n,\Phi_n}(\Upsilon_n)}(\Upsilon_n) > \mu_{\Upsilon_n} + \frac{\delta}{2}\right) + \Pr\left(\mathcal{K}(\mu_{\Upsilon_n} + \frac{\delta}{2}, \mu_{\Upsilon_n} + \delta) \leq \frac{\mathcal{F}(\Phi_n)}{N'_{\Gamma_n,\Phi_n}(\Upsilon_n)}\right) \\
&\leq \sum_{a\in\mathcal{A}}\sum_{p=1}^{M} \Pr\left(\hat{\mu}_{p,N_{\Gamma_n,\Phi_n}(\Upsilon_n)}(a) < \mu_a - \delta\right) \qquad \text{(G.2)} \\
&+ \sum_{a\in\mathcal{A}}\sum_{p=1}^{M} \Pr\left(\hat{\mu}_{p,N_{\Gamma_n,\Phi_n}(\Upsilon_n)}(a) > \mu_a + \frac{\delta}{2}\right) \qquad \text{(G.3)} \\
&+ \Pr\left(\min_{a\in\mathcal{A}} \mathcal{K}(\mu_a + \frac{\delta}{2}, \mu_a + \delta) \leq \frac{\mathcal{F}(\Phi_n)}{N'_{\Gamma_n,\Phi_n}(\Upsilon_n)}\right). \qquad \text{(G.4)}
\end{aligned}$$

Recall that $N_{\Gamma_n,\Phi_n}(\Upsilon_n) \geq (\Phi_n)^\xi$ for n large enough, and $\Phi_n \xrightarrow{p} \infty$. As a consequence,

$$N_{\Gamma_n,\Phi_n}(\Upsilon_n) \xrightarrow{p} \infty \text{ as } n \to \infty.$$

Then by Lemma 21, both (G.2) and (G.3) converge to 0 as $n \to \infty$. For (G.4), on the one hand, $\min_{a\in\mathcal{A}} \mathcal{K}(\mu_a + \frac{\delta}{2}, \mu_a + \delta)$ is a constant; on the other hand, $N'_{\Gamma_n,\Phi_n}(\Upsilon_n) \geq N_{\Gamma_n,\Psi_n}(\Upsilon_n) \geq (\Phi_n)^\xi$ for sufficiently large n and $(\Phi_n)^\xi \in \omega\left(\mathcal{F}(\Phi_n)\right)$, therefore

$$\frac{\mathcal{F}(\Phi_n)}{N'_{\Gamma_n,\Phi_n}(\Upsilon_n)} \to 0 \text{ as } n \to \infty.$$

Hence (G.4) is always 0 for n large enough. This concludes the proof. □

Lemma 23. *Let $(\Phi_n)_{n\geq 1}$ be a sequence of random variables such that $\Phi_n \xrightarrow{p} \infty$. Then for every player p and every arm a,*

$$\tilde{\mu}_{p,\Phi_n}(a) \xrightarrow{p} \mu_a.$$

Proof. First note that $\tilde{\mu}_{p,\Phi_n}(a) = \hat{\mu}_{p,N_{p,\Phi_n}(a)}(a)$. By Lemma 21 it suffices to prove that $N_{p,\Phi_n}(a) \xrightarrow{p} \infty$. By the definition of convergence in probability, we only need to show that for every $N > 0$ and every $\epsilon > 0$, we can find an N_* such that for any $n \geq N_*$,

$$\Pr\left(N_{p,\Phi_n}(a) < N\right) \leq \epsilon.$$

For each $n \geq 1$, let Υ_n be the arm that player p has pulled most frequently by round Φ_n and let Ψ_n be the last round in the first Φ_n rounds that p pulled Υ_n. The definition of Ψ_n implies that

$$\Psi_n \geq \Phi_n \,/\, K.$$

Therefore given $\Phi_n \xrightarrow{p} \infty$ we have $\Psi_n \xrightarrow{p} \infty$. In addition, we have

$$N_{p,\Psi_n}(\Upsilon_t) \geq \Phi_n \,/\, K \geq \Psi_n \,/\, K.$$

Hence by Lemma 22,

$$\lim_{t\to\infty} \Pr\left(B^+_{p,\Psi_n}(\Upsilon_n) > \mu_{\Upsilon_n} + \epsilon\right) = 0, \text{ for any } \epsilon > 0. \qquad \text{(G.5)}$$

Now let $\delta = (1 - \max_{a\in\mathcal{A}} \mu_a) \,/\, 2$. Then for every $N > 0$ we have

$$\begin{aligned}
\Pr\left(N_{p,\Phi_n}(a) < N\right) &= \Pr\left(N_{p,\Phi_n}(a) < N \wedge B^+_{p,\Psi_n}(\Upsilon_n) \geq B^+_{p,\Psi_n}(a)\right) \\
&\leq \Pr\left(B^+_{p,\Psi_n}(\Upsilon_n) > 1 - \delta\right) + \Pr\left(\mathcal{K}'(\tilde{\mu}_{p,\Psi_n}(a), 1 - \delta) > \frac{\mathcal{F}(\Psi_n)}{N \cdot M}\right) \\
&\leq \Pr\left(B^+_{p,\Psi_n}(\Upsilon_n) > \mu_{\Upsilon_n} + \delta\right) + \Pr\left(N \cdot M \cdot \mathcal{K}(0, 1 - \delta) > \mathcal{F}\left(\frac{\Phi_n}{K}\right)\right).
\end{aligned}$$

For every $\epsilon > 0$, by (G.5) we can choose N_0 large enough such that for every $n \geq N_0$, $\Pr(B^+_{p,\Psi_n}(\Upsilon_n) > \mu_{\Upsilon_n} + \delta) < \epsilon \,/\, 2$. Since $N \cdot M \cdot \mathcal{K}(0, 1-\delta)$ is a constant, we can also choose N_1 large enough such that for every $n \geq N_1$, $N \cdot M \cdot \mathcal{K}(0, 1-\delta) \leq \mathcal{F}(n \,/\, K)$. Finally, according to the assumption that $\Phi_n \xrightarrow{p} \infty$, we can choose N_2 large enough such that for every $n \geq N_2$, $\Pr(\Phi_n < N_2) \leq \epsilon \,/\, 2$. Thus, for every $n \geq N_* = \max(N_0, N_1, N_2)$, $\Pr\left(N_{p,\Phi_n}(a) < N\right) < \epsilon$. $\square$

Then we can start our proof of Lemma 20. In fact, it is equivalent to prove the following equation:

$$\lim_{n\to\infty} \Pr\left(N_{\Phi_n}(a) < \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})}\right) = 0.$$

For every round t, we define a random player p_t such that

$$N'_{p_t,t}(a) \leq N'_{p',t}(a), \quad \text{for every } p' \in [M] \tag{G.6}$$

and we say player p_t is *chosen* for round t. For every n, we define a random player Γ_n such that Γ_n is the player who is chosen most frequently in the first Φ_n rounds. We also let Ξ be a random set including the rounds (in the first T rounds) in which player Γ_n is chosen. We let random variable Υ_n be the arm that is pulled most frequently by Γ_n in those rounds in Ξ_n and let Ψ_n be the last round in Ξ_n that Γ_n pulls Υ_n. With all the definitions above, we have

$$N_{\Gamma_n,\Psi_n}(\Upsilon_n) \geq \frac{\Phi_n}{MK} \geq \frac{\Psi_n}{MK}. \tag{G.7}$$

as well as

$$\Psi_n \geq \frac{\Phi_n}{MK}. \tag{G.8}$$

Note that $\Phi_n \xrightarrow{p} \infty$ implies $\Phi_n \xrightarrow{p} \infty$, which by (G.8) in turn implies

$$\Psi_n \xrightarrow{p} \infty \tag{G.9}$$

For any fixed T, clearly we have

$$\begin{aligned} N_{\Phi_n}(a) < \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})} \implies \\ N_t(a) < \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})}, \quad \text{for every } t \leq \Phi_n. \end{aligned}$$

By Claim 11 and (G.6),

$$\begin{aligned} N_{\Phi_n}(a) < \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})} \implies \\ N'_{p_t,t}(a) \leq \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})}, \quad \text{for every } t \leq \Phi_n. \end{aligned}$$

Furthermore, by the definitions of random variables Γ_n, Ψ_n, and Υ_n,

$$N_{\Phi_n}(a) < \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})} \implies \begin{cases} N'_{\Gamma_n,\Psi_n}(a) < \dfrac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})} \\ B^+_{\Gamma_n,\Psi_n}(\Upsilon_n) \geq B^+_{\Gamma_n,\Psi_n}(a) \end{cases}$$

Hence we have

$$\begin{aligned} &\lim_{n\to\infty} \Pr\left(N_{\Phi_n}(a) < \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})}\right) \\ &\quad \leq \lim_{n\to\infty} \Pr\left(N'_{\Gamma_n,\Psi_n}(a) < \frac{(1-\delta)\mathcal{F}(\Phi_n)}{\mathcal{K}(\nu_a, \nu_{a^*})} \wedge B^+_{\Gamma_n,\Psi_n}(\Upsilon_n) \geq B^+_{\Gamma_n,\Psi_n}(a)\right) \\ &\quad\quad \leq \lim_{n\to\infty} \Pr\left(B^+_{\Gamma_n,\Psi_n}(\Upsilon_n) > \mu_{\Upsilon_n} + \epsilon\right) &\text{(G.10)} \\ &\quad\quad + \sum_{p=1}^{M} \lim_{n\to\infty} \Pr\left(\tilde{\mu}_{p,\Psi_n}(a) < \mu_a - \epsilon\right) &\text{(G.11)} \\ &\quad\quad + \lim_{n\to\infty} \Pr\left(\mathcal{K}(\mu_a - \epsilon, \mu^* + \epsilon) \geq \frac{\mathcal{F}(\Psi_n)\mathcal{K}(\nu_a, \nu_{a^*})}{(1-\delta)\mathcal{F}(\Phi_n)}\right), &\text{(G.12)} \end{aligned}$$

where ϵ is to be determined later. By Lemma 22 and (G.7), the first term (G.10) is 0. By Lemma 23 and (G.9), the second term (G.11) is 0. For the third term (G.12), by (G.8) we have

$$\begin{aligned}
&\lim_{n\to\infty} \Pr\left(\mathcal{K}(\mu_a - \epsilon, \mu^* + \epsilon) \geq \frac{\mathcal{F}(\Psi_n)\mathcal{K}(\nu_a, \nu_{a^*})}{(1-\delta)\mathcal{F}(\Phi_n)}\right) \\
&\quad \leq \lim_{n\to\infty} \Pr\left(\mathcal{K}(\mu_a - \epsilon, \mu^* + \epsilon) \geq \frac{\mathcal{F}\left(\frac{\Phi_n}{MK}\right)\mathcal{K}(\nu_a, \nu_{a^*})}{(1-\delta)\mathcal{F}(\Phi_n)}\right)
\end{aligned}$$

Now we decide ϵ to be a positive real number small enough such that there exists a T_0 satisfying

$$\mathcal{K}(\mu_a - \epsilon, \mu^* + \epsilon) < \frac{\mathcal{F}(t \,/\, (MK))\mathcal{K}(\nu_a, \nu_{a^*})}{(1-\delta)\ln(t)}, \quad \text{for every } t \geq T_0.$$

Thus,

$$\lim_{n\to\infty} \Pr\left(\mathcal{K}(\mu_a - \epsilon, \mu^* + \epsilon) \geq \frac{\mathcal{F}\left(\Phi_n \,/\, (MK)\right)\mathcal{K}(\nu_a, \nu_{a^*})}{(1-\delta)\ln(\Phi_t)}\right) \leq \lim_{n\to\infty} \Pr\left(\Phi_n < T_0\right) = 0.$$

This concludes the proof.

Appendix H Proof of Claim 11

It suffices to prove $\sum_{p=1}^{M} N'_{p,t}(a) \leq M \cdot N_t(a)$. In fact,

$$\begin{aligned}
\sum\nolimits_{p=1}^{M} N'_{p,t}(a) &\leq \sum\nolimits_{p=1}^{M} \left(N_{p,t}(a) + (M-1)\cdot\left(N_{p,t}(a) - N_{\ell(t)}(a)\right)\right) \qquad \text{(by Definition)} \\
&\leq \sum\nolimits_{p=1}^{M} \left(N_{\ell(t)}(a) + M\cdot\left(N_{p,t}(a) - N_{\ell(t)}(a)\right)\right) \\
&\leq M\cdot N_{\ell(t)}(a) + M\cdot\sum\nolimits_{p=1}^{M}(N_{p,t}(a) - N_{\ell(t)}(a)) \\
&\leq M\cdot N_{\ell(t)}(a) + M\cdot(N_t(a) - N_{\ell(t)}(a)) \\
&= M\cdot N_t(a),
\end{aligned}$$

which concludes the proof.

Appendix I Proof of Claim 12

If $\alpha(\mathcal{C}) = 0$, then the right hand side becomes $M \cdot N_{p,t}(a)$. By definition,

$$\begin{aligned}
N'_{p,t}(a) &\leq N_{p,t}(a) + (M-1)\cdot\left(N_{p,t}(a) - N_{\ell(t)}(a)\right) \\
&\leq N_{p,t}(a) + (M-1)\cdot N_{p,t}(a) \\
&= M\cdot N_{p,t}(a).
\end{aligned}$$

If $N_{p,t}(a) = 0$, then $N'_{p,t}(a) = 0$ and the bound is trivial. Now we assume $\alpha(\mathcal{C}) > 0$ and $N_{p,t}(a) > 0$. Note that

$$\frac{N'_{p,t}(a)}{N_{p,t}(a)} = 1 + (M-1)\cdot\min\left(1 - \frac{N_{\ell(t)}(a)}{N_{p,t}(a)}, \frac{u(\ell(t))}{N_{p,t}(a)}\right).$$

Let $f(x) = \min\left(1 - N_{\ell(t)}(a) \,/\, x, u(\ell(t)) \,/\, x\right)$. We have

$$\frac{N'_{p,t}(a)}{N_{p,t}(a)} \leq 1 + (M-1)\cdot\left(\sup_{x\in(0,\infty)} f(x)\right).$$

Since $1 - N_{\ell(t)}(a) \,/\, x$ is increasing in $(0,\infty)$ and $u(\ell(t)) \,/\, x$ is decreasing in $(0,\infty)$, $f(x)$ can be maximized if these two functions take the same value. In fact, when $x = x^* = N_{\ell(t)}(a) + u(\ell(t))$,

we have $f(x^*) = 1 - N_{\ell(t)}(a) \,/\, x^* = u(\ell(t)) \,/\, x^*$. Thus,

$$\begin{aligned}
\frac{N'_{p,t}(a)}{N_{p,t}(a)} &\leq 1 + (M-1) \cdot \frac{u(\ell(t))}{N_{\ell(t)}(a) + u(\ell(t))} \\
&\leq 1 + (M-1) \cdot \frac{\frac{N_{\ell(t)}(a)/\alpha(\mathcal{C}) - N_{\ell(t)}(a)}{M}}{N_{\ell(t)}(a) + \frac{N_{\ell(t)}(a)/\alpha(\mathcal{C}) - N_{\ell(t)}(a)}{M}} \\
&= \frac{M}{1 + (M-1)\alpha(\mathcal{C})},
\end{aligned}$$

which completes the proof.